Bundle Pricing with Comparable Items

Alexander Grigoriev[1], Joyce van Loon[1,⋆], Maxim Sviridenko[2], Marc Uetz[1], and Tjark Vredeveld[1]

[1] Maastricht University, Quantitative Economics,
P.O. Box 616, NL–6200 MD Maastricht, The Netherlands
{a.grigoriev,j.vanloon,m.uetz,t.vredeveld}@ke.unimaas.nl
[2] IBM T.J. Watson Research Center,
P.O. Box 218, Yorktown Heigths, NY 10598, USA
sviri@us.ibm.com

Abstract. We consider a revenue maximization problem where we are selling a set of items, each available in a certain quantity, to a set of bidders. Each bidder is interested in one or several bundles of items. We assume the bidders' valuations for each of these bundles to be known. Whenever bundle prices are determined by the sum of single item prices, this algorithmic problem was recently shown to be inapproximable to within a semi-logarithmic factor. We consider two scenarios for determining bundle prices that allow to break this inapproximability barrier. Both scenarios are motivated by problems where items are different, yet comparable. First, we consider classical single item prices with an additional monotonicity constraint, enforcing that larger bundles are at least as expensive as smaller ones. We show that the problem remains strongly NP-hard, and we derive a PTAS. Second, motivated by real-life cases, we introduce the notion of affine price functions, and derive fixed-parameter polynomial time algorithms.

1 Introduction

Consider the situation that we want to sell a set of items to a set of bidders. Every bidder places bids on subsets, or *bundles* of items, and each bidder would like to receive one or more of these bundles (OR-bids). Bidders have valuations for each of their bundles. The valuation is the maximum amount a bidder is willing to pay for a particular bundle. We assume that the bundles and the valuations are known. Hence we are faced with a purely algorithmic problem, in contrast to a mechanism design problem where the valuations are private information to the bidders. We have a certain amount of copies of each item available, and this amount may be limited or unlimited, for example non-digital or digital goods. We need to determine two things, namely which of the bidders receive which of their requested bundles, and how much each of them needs to pay. The goal is to maximize the total revenue received from the bidders. A

⋆ Supported by METEOR, the Maastricht Research School of Economics of Technology and Organizations.

L. Arge, M. Hoffmann, and E. Welzl (Eds.): ESA 2007, LNCS 4698, pp. 475–486, 2007.

general economic constraint on the possible prices, adopted in this paper as well, is that of *envy-freeness*. It requires that no bidder is left envious in the sense that she could afford a bundle, but doesn't receive it[1]. This is the general setting for the pricing problems studied in this paper.

In a sequence of recent papers [1,2,4,5,8,9,10], several algorithms and complexity results have been derived for such price optimization problems. The pricing model that is assumed in all these papers is the problem with *single item prices*, where each item is assigned an anonymous price, and bundle prices are defined by the sum of the respective item prices. We contribute to this line of research in two different directions.

First, we consider the single item prices model. We introduce a monotonicity constraint that allows us to derive results that break the semi-logarithmic inapproximability barrier known for the general case [5]. We impose the condition that the price of any bundle of size k must not exceed the price of a set of size $k+1$ or larger, for any k. This condition implies that (most of) the items for sale are *comparable* in the sense that the prices do not differ too much.

Second, we propose a model for determining bundle prices that actually generalizes the single item prices problem. We derive fixed-parameter polynomial time algorithms for that model. We assume that the bundle prices are determined on the basis of arbitrary affine functions, defined on a joint set of variables. Let us give an illustrating example: For each bundle j, there is an individual fixed cost a_j that the bidder needs to pay when purchasing that bundle. All items are identical and the seller just needs to determine one per-item price, say x. The price for any bundle j of size k_j is then given by $a_j + k_j x$. Notice that the price paid for any bundle is an affine function that depends on the size of the requested bundle. In general, the affine pricing model allows for many more pricing scenarios relevant in practice, e.g. groups of customers with different price functions, quantity discounts, etc. See Section 3 for a brief discussion.

Before we elaborate on related work and our contribution, let us define the pricing settings more formally.

1.1 Model

Let $I = \{1, \ldots, m\}$ denote the set of comparable items for sale, and let $J = \{1, \ldots, n\}$ denote the set of bids placed by all bidders. Each bid $j \in J$ is on exactly one subset of items $I_j \subseteq I$. In line with notation in auction literature, we call the set I_j also a *bundle*. Every bidder has a positive valuation for each of her bundles, that is, every bundle I_j corresponding to bid $j \in J$, has a positive valuation b_j which is the maximum amount its bidder is willing to pay for bundle I_j. We may assume w.l.o.g. that $b_j \geq 1$ for $j \in J$. The valuations are assumed to be known to the seller. Let c_i denote the available number of copies of item $i \in I$. We consider both the case of unlimited availability of items, that is, $c_i \geq n$ for all items $i \in I$, and the case of limited availability of items.

[1] More generally, envy-freeness requires that in an allocation, the bundle allocated to a bidder belongs to her *demand set*, which is the set of all allocations that maximize the bidder's utility [11].

A bid is a *winning bid* if it is assigned to the bidder, and a *losing bid* otherwise. The set of winning bids is denoted by $W \subseteq J$. A solution to the problem is a price $p(j)$ for the bundle I_j corresponding to bid $j \in J$. Later we will be more specific about further restrictions on the prices. A solution is called *feasible* if the bundles of all winning bids can be afforded by the respective bidder (that is, the price of the bundle corresponding to a bid is at most its valuation), and if no item is oversold. A solution is *envy-free* if in addition, for all losing bids the respective bundle is priced higher than the valuation of the corresponding bid. Let us summarize the above discussion in a definition for the generic pricing problem that we address in the paper.

Definition 1. *A feasible and envy-free solution to a pricing problem consists of a price $p(j)$ for the bundle I_j, for all bids $j \in J$, and a set of winning bids $W \subseteq J$ that are assigned to their bidders such that*

1. *the bundle of every winning bid j is affordable for the bidder, that is $p(j) \leq b_j$ for all $j \in W$,*
2. *the bundle of every losing bid j is too expensive for the bidder, that is $p(j) > b_j$ for all $j \in J \setminus W$,*
3. *no item is oversold, that is $\sum_{j \in W} |\{i\} \cap I_j| \leq c_i$ for all items $i \in I$.*

The objective is to find a solution that maximizes the total revenue of the seller, that is, we want to maximize $\sum_{j \in W} p(j)$.

We consider two different models for the computation of prices. The first pricing model that we consider is the single item prices model, where we have to determine item prices for all items $i \in I$. To be in line with previous papers on the same topic, let p_i denote the price of item i, and the price of bundle $I_j \subseteq I$ is $p(j) = \sum_{i \in I_j} p_i$, for all $j \in J$. Given that several inapproximability results exist for this model [5,8,9], we introduce a monotonicity constraint on the set of item prices. Specifically, we impose that the following holds true for any two subsets of items I' and I''.

$$p(I') \leq p(I'') \text{ whenever } |I'| < |I''| \,. \tag{1}$$

The condition has a meaningful economic interpretation in a lot of settings where items are different yet comparable, as it only requires that larger bundles are at least as expensive as smaller ones. It yields that (most) item prices are of the same order of magnitude. We show how this monotonicity constraint can be exploited to derive results that break the inapproximability barrier known for the general unconstrained case.

In practice, bundle prices are often not determined by the sum of individual item prices, but rather by a function based on a few bundle characteristics. Therefore, we propose the second model, in fact generalizing the single item prices problem, see Proposition 1. Here, the price of bundle I_j is determined by an affine function in some dimension K as follows.

$$p(j) = a_{j0} + a_{j1}x_1 + \cdots + a_{jK}x_K \,, \qquad j \in J \,. \tag{2}$$

The coefficients a_{jk}, $k = 0, \ldots, K$, are arbitrary coefficients that are given for all bids $j \in J$. These coefficients may, in general, depend on both the bundle I_j and on the bidder that places bid j. Thus it may be the case that similar bundles have different prices. The pricing problem consists of determining revenue-maximizing values for the (nonnegative) variables x_k, $k = 1, \ldots, K$. We postpone the discussion of this model to Section 3.

In the remainder of this paper, we denote by a ρ-approximation algorithm an algorithm that produces a solution with value at least $1/\rho$ times the optimal solution value. A PTAS is a family of $(1+\varepsilon)$-approximation algorithms, for any $\varepsilon > 0$.

1.2 Related Work

The problem mainly addressed in the literature is the one with unlimited availability of items, single item prices, and the requirement that the solution is envy-free [1,2,4,5,9,10]. For this problem the maximum revenue is hard to approximate to within a semi-logarithmic factor in the number of bids n [5]. In particular, it is unlikely that a constant approximation algorithm exists. For the same problem, Hartline and Koltun [10] present an approximation scheme with almost linear running time, given that the number of distinct items is constant. Moreover, given that each bidder is interested in bundles of at most k items, Balcan and Blum [2] derive an $\mathrm{O}(k)$-approximation. Finally, there exist two fully polynomial time approximation schemes [2,4] for the problem where the bundles are *nested*, that is, for any two bundles I_j and $I_{j'}$ it holds that $I_j \subseteq I_{j'}$, $I_{j'} \subseteq I_j$ or $I_j \cap I_{j'} = \emptyset$.

1.3 Our Results

For the revenue maximization problem with single item prices, we derive strong NP-hardness even if prices need to fulfill the monotonicity constraint. Moreover, we derive a PTAS for that problem, with a time complexity of $\mathrm{O}(nm^{8/\varepsilon}(\log B))^{8/\varepsilon})$, where $B = \max_j b_j$.

For the revenue maximization problem with affine price functions, we propose an algorithm with a time complexity of $\mathrm{O}((K^3+nK)(n+K)^K)$. Here, parameter K is the dimension in which the affine price functions live. In particular, for $K = 1$ this is $\mathrm{O}(n^2)$. A similar result (with slightly different time complexity) holds for the case of limited availability of items. For the same problem with non-constant K, and unlimited availability of items, the maximum revenue is hard to approximate to within a semi-logarithmic factor in the number of bids n. This follows directly from the corresponding result by Demaine et al. [5], as we can show that the problem with single item prices is a special case. In addition, for the same pricing problem with limited availability of items, we prove that it is even NP-complete to approximate the maximum revenue to within a factor of $n^{1-\varepsilon}$ of optimum, where n is the number of bids.

A special case of single item pricing is the so-called highway pricing problem as suggested in [9]. There the bundles are subpaths of a simple path. We show that this problem remains NP-hard even under the monotonicity assumption, and we derive a simple $O(\log B)$-approximation algorithm, where $B = \max_{j\in J} b_j$.

2 Single Item Pricing with Monotonicity Constraint

In single item pricing, we need to assign an (anonymous) item price p_i for each item $i \in I$, and bundle prices $p(j)$ are defined as the sum of the prices of the requested items, $p(j) = \sum_{i\in I_j} p_i$, for all $j \in J$. As before, the item prices need to yield a feasible and envy-free solution, and we wish to maximize the total revenue, which can be written as $\sum_{j\in W} \sum_{i\in I_j} p_i$. Notice that in case of unlimited availability of items both feasibility and envy-freeness is in fact no issue – yet finding optimal prices is hard [5]. For this reason we introduce a *monotonicity constraint*: $p(I') \leq p(I'')$ if $|I'| < |I''|$, for any two subsets of items I' and I''.

2.1 Complexity

Theorem 1. *The revenue maximization problem with single item prices and unlimited availability of items is strongly NP-hard, even if the prices need to satisfy the monotonicity constraint.*

Proof. We use a reduction from the strongly NP-hard problem IndependetSet [6]. Let $G = (V, E)$ be a graph in which we want to find a maximum cardinality set of vertices that are pairwise non-adjacent. Let M be an integer that is large enough. For every vertex $v \in V$ we create a *vertex-item*, and for every edge $e \in E$ we introduce an *edge-item*, that is, $I = V \cup E$. For every item $i \in I$, there are $M + 2$ bids placed on the bundle consisting of only this item. One of these bids has valuation M, and the others have the same valuation $M + 1$. Moreover, for every edge $e = \{u, v\} \in E$, there are four more bids. One bid is on bundle $\{u, e\}$, one bid on bundle $\{v, e\}$, and two bids are on bundle $\{u, v\}$. These four bids each have valuation $2M + 1$.

We claim that there exists an independent set of size s in G if and only if there is a solution for the revenue maximization problem with revenue $f + s$, where f is a function of M, $|V|$ and $|E|$. To prove this claim, we show that in any optimal solution all items are priced either at M or at $M + 1$. Moreover, vertex-items are priced at $M + 1$ if and only if the corresponding vertex belongs to the independent set. Details of the proof are included in the full version of this paper. $\square$

2.2 Approximation Scheme

In order to derive a PTAS for the problem with single item prices and monotonicity constraint, we restrict the prices to powers of $(1 + \delta)$ for some $\delta > 0$. Assume, without loss of generality, that $p_1 \leq p_2 \leq \ldots \leq p_m$, then by the monotonicity constraint, we know $2p_2 \geq p_1 + p_2 \geq p_m$. Similarly, $3p_3 \geq p_1 + p_2 + p_3 \geq p_{m-1} + p_m \geq 2p_{m-1}$, etc.

Lemma 1. *Suppose $p_1 \leq p_2 \leq \ldots \leq p_m$. Any pricing of the items satisfying the monotonicity constraint also satisfies*

$$k\, p_k \geq (k-1) p_{m-k+2}, \quad k = 2, \ldots, \left\lceil \frac{m}{2} \right\rceil . \tag{3}$$

The idea for the PTAS is now the following. Except for a constant number of the cheapest and most expensive items, all items have prices in roughly the same range. Therefore we can price all except a constant number of items uniformly with the same price, without loosing too much in terms of the total revenue. We therefore enumerate over all possible uniform prices for the bulk of the items, and over all possible combinations of prices for the remaining (constant number of) items.

Theorem 2. *The pricing problem with unlimited availability of items, single item prices and monotonicity constraint admits a PTAS. The time complexity is $O(nm^{8/\varepsilon}(\log B)^{8/\varepsilon})$, where ε is the precision of the PTAS and $B = \max_j b_j$.*

Proof. Given an instance of the pricing problem and an $\varepsilon > 0$, let $\delta = \varepsilon/4$, and for convenience assume that $1/\delta$ is integral. Assume that we know the order of prices, say $p_1 \leq \cdots \leq p_m$, in an optimum solution. Define the subsets of items $S = \{i \in I : i \leq \frac{1}{\delta}\}$, $M = \{i \in I : 1 + \frac{1}{\delta} \leq i \leq m + 1 - \frac{1}{\delta}\}$ and $L = \{i \in I : i \geq m + 2 - \frac{1}{\delta}\}$. Note that $M = \emptyset$ if $\varepsilon \leq 8/(m+1)$, in which case the number of items is in $O(1/\varepsilon)$. We round down the prices of all items in S and L to powers of $(1+\delta)$. Moreover, we price all items in M uniformly at price $p_{1+1/\delta}$, rounded down to a power of $(1+\delta)$. Let us call the new prices p', and let us call p'_M the price of items in M. First observe that the order of prices does not change. We next argue that we do not loose too much by this rounding. Clearly, since we only round down, the set of winning bids can only increase. Moreover, we loose at most a factor $(1+\delta)$ on items in S and L. Finally, consider the items in M. By (3), we have

$$\left(1 + \frac{1}{\delta}\right) p_{1+1/\delta} \geq \frac{1}{\delta} p_{m+1-1/\delta} \, .$$

In other words, the price for the most expensive item in M differs from the cheapest item in M by a factor at most $(1+\delta)$. Hence, on items in M we loose a factor at most $(1+\delta)^2$.

Now we have a structured solution, but it may violate the monotonicity constraint. We claim that any such violation can be restored by one more rounding operation, if necessary: We just round down the price of all items priced p'_M or higher by another factor $(1+\delta)$. For contradiction, after this last rounding consider two violating sets I' and I'' with $|I'| < |I''|$ and $p'(I') > p'(I'')$, and w.l.o.g. $|I'| = \ell$, and $|I''| = \ell + 1$. Due to the ordering of prices, we then also have that $p'(\{1, \ldots, \ell+1\}) < p'(\{m, m-1, \ldots, m-\ell+1\})$. As long as there are items from M in both sets, we redefine $\ell = \ell - 1$, and we keep violating the monotonicity constraint. But now, all items in $\{1, \ldots, \ell\}$ have been rounded down by a factor at most $(1+\delta)$, and all items in $\{m, m-1, \ldots, m-\ell+1\}$ have been rounded down by a factor at least $(1+\delta)$. This contradicts the monotonicity constraint of the optimal solution that we started with.

The PTAS now consists of enumerating all possible structured solutions, which is sufficient to obtain a feasible solution that differs from the optimal solution by a factor at most $(1+\delta)^3 < (1+\varepsilon)$. There are $\binom{m}{-1+2/\delta}$ possible choices for $S \cup L$. Since all prices are powers of $(1+\delta)$, there are $\log B$ possible prices. Given that all items in M have the same price, there are at most $(\log B)^{2/\delta}$ structured solutions for each choice of $S \cup L$. Computation of the revenue for any such solution takes $\mathrm{O}(nm)$ time. This together with $\delta = \varepsilon/4$ yields the claimed time complexity, where the constant hidden in the O-notation depends on ε. □

3 Pricing with Affine Price Functions

For this section, we address revenue maximization problems with bundle prices $p(j)$ that are determined via affine price functions $p(j) = a_{j0} + a_{j1}x_1 + \cdots + a_{jK}x_K$, one affine function per bid $j \in J$, as defined in (2).

Let us discuss examples to motivate this model. If we let $K = 1$ and define $a_{j1} = |I_j|$ for all bids j, the bundle prices are determined by affine functions that depend only on the size of the bundles. The optimization problem is to determine the per-item price x_1, which is identical for all items. Fixed costs per bid j are incorporated by letting $a_{j0} \neq 0$. Bidder-dependent characteristics are easily incorporated as well, for example cost reductions of $\alpha_t\%$ for certain types t of customers, e.g. letting $a_{j2} = -\alpha_t a_{j1}$. Another meaningful interpretation is this: There are K different item types $i = 1, \ldots, K$ for which we need to determine the per-item prices x_i, and any bid j is specified by the number of requested items of type i, a_{ji}, and the total valuation b_j. In fact, one motivation for this model is phone contracts, where x_1 represents the monthly subscription fee ($a_{j1} = 1$ for all j), x_2 is the price per SMS and x_3 is the price per minute for phone calls. Here, the coefficients a_{j2} and a_{j3} describe typical average usages for different types j of customers.

We distinguish between unlimited and limited availability of items.

3.1 Unlimited Availability of Items

Our algorithm is polynomial as long as K, the dimension of the affine price functions is constant. It simply enumerates all vertices of the linear arrangement defined by the *valuation constraints*. Here, the valuation constraints are given by $p(j) \leq b_j$ for every winning bid j. More precisely, suppose that we know which of the bids are winning bids in an optimum solution, say $W \subseteq J$. Then we know that the variables $x_1, \ldots, x_K$ have to fulfill the $|W|$ inequalities

$$a_{j0} + a_{j1}x_1 + \cdots + a_{jK}x_K \leq b_j\,, \quad j \in W\,.$$

Denote by $\mathcal{P}$ the polyhedron defined by these $|W|$ inequalities. For an optimum solution $x = (x_1, \ldots, x_K)$, at least one of these inequalities must be tight, because otherwise the bundles corresponding to the same set of winning bids could be priced even higher. Assume that $W' \subseteq W$ are the bids for which the above inequalities are tight, and note that W' is nonempty. Then the system

$$a_{j0} + a_{j1}x_1 + \cdots + a_{jK}x_K = b_j \,, \quad j \in W'$$

defines a (nonempty) face $\mathcal{F}$ of the polyhedron $\mathcal{P}$. By definition, any point $x \in \mathcal{F}$ defines an optimal solution. Clearly, at most K inequalities are required to completely characterize the face $\mathcal{F}$. Moreover, we have exactly $\dim(\mathcal{F})$ free variables in the optimal solution x. In other words, the same total revenue can be obtained by fixing the $\dim(\mathcal{F})$ free variables among $x_1, \ldots, x_K$ to 0. Hence, an optimal solution can be obtained by considering all solutions x that are characterized by K linearly independent constraints out of the following $n + K$ constraints.

$$a_{j0} + a_{j1}x_1 + \cdots + a_{jK}x_K = b_j \,, \qquad j \in J \,, \tag{4}$$

$$x_k = 0 \,, \qquad k = 1, \ldots, K \,. \tag{5}$$

This insight can be used to define a simple algorithm that solves the revenue maximization problem in polynomial time, as long as K is constant.

Algorithm 1. Revenue maximization with affine price functions

Input: Instance with affine price functions as defined in (2).

For all candidate solutions $x = (x_1, \ldots, x_K)$ that fulfill K linearly independent constraints out of the $n + K$ constraints (4) and (5) do:
–Let $p(j) = a_{j0} + a_{j1}x_1 + \cdots + a_{jK}x_K$, $j \in J$, be the bundle prices.
–Let $W := \{j \in J \,:\, p(j) \leq b_j\}$ be the set of winning bids.
–Let $\Pi = \sum_{j \in W} p(j)$ be the total revenue.

Output: Maximum among all values Π, with optimal parameters $x_1, \ldots, x_K$, and set of winning bids W.

Theorem 3. *Algorithm 1 solves the revenue maximization problem with affine price functions and unlimited availability of items in $O((K^3 + nK)(n + K)^K)$ time.*

Proof. Correctness of the algorithm immediately follows from the preceding discussion. We need to consider $\binom{n+K}{K} \in \mathrm{O}((n + K)^K)$ systems of K constraints each. In each of these iterations, we need to solve a linear system in K variables and K constraints, which takes $\mathrm{O}(K^3)$ time. Computation of the bundle prices, winning bids, and the objective value takes $\mathrm{O}(nK)$ time. The claimed time complexity follows. □

In contrast, if the dimension K of the price functions is not constant, the problem is much harder. In fact, if K is not constant, the single item prices problem is just a special case of the problem with affine price functions: Let $K = m$, and let $a_{ji} = 1$ whenever item i is contained in bundle I_j and $a_{ji} = 0$ otherwise. Then, each item price corresponds to one variable x_i, and we immediately obtain the following.

Proposition 1. *The model with affine price functions generalizes the model with single item prices. Hence, the semi-logarithmic inapproximability result by Demaine et al. [5] also holds for pricing problems with affine price functions.*

3.2 Limited Availability of Items

First we claim that Algorithm 1 can as well be used to solve the problem when the availability of items is limited. Indeed, the only thing we additionally need to check for any of the candidate solutions is feasibility: We have to verify whether none of the items is oversold, and if yes, we do not consider the candidate solution. Also note that any candidate solution is envy-free by definition. Clearly, this feasibility check can be done in $O(nm)$ time per candidate solution.

Corollary 1. *Algorithm 1, augmented with a feasibility check, solves the revenue maximization problem with affine price functions and limited availability of items in $O((K^3 + nK + nm)(n + K)^K)$ time.*

On the negative side, it turns out that the problem with limited availability of the items seems even harder to approximate.

Theorem 4. *Consider the revenue maximization problem with affine price functions and limited availability of items. For any $\varepsilon > 0$, it is NP-hard to approximate the maximum revenue to within a factor $n^{1-\varepsilon}$. This result holds even if all bids have a valuation of one, the availability of each item is one, and each item is requested in at most two bids.*

Proof. We use an approximation preserving reduction from IndependetSet. Given is a graph $G = (V, E)$, we want to find a maximum cardinality subset $V' \subseteq V$ such that no two vertices in V' are adjacent. Zuckerman [12] showed that it is NP-hard to approximate the maximum independent set to within a factor $|V|^{1-\varepsilon}$, for any $\varepsilon > 0$.

We construct the following instance of the pricing problem. Each vertex $v \in V$ corresponds to a bid and each edge $e \in E$ corresponds to an item. Each bid v is on a bundle containing all edges incident to v, and has valuation $b_v = 1$. Each item e is available once ($c_e = 1$). Let the price functions be $p(v) = 1 + x_v$, for all bundles I_v corresponding to bid $v \in V$.

We claim that an independent set of cardinality s exists in G if and only if there exists a pricing for the above defined instance with total revenue s. Suppose $V' \subseteq V$ is an independent set in G, $|V'| = s$. Then let $x_v = 0$ for all $v \in V'$, and $x_v > 0$ otherwise. This way the set of winning bids equals the independent set V', and therefore no item is oversold. No bidder is envious, as the price of a bundle corresponding to a losing bid exceeds its valuation, and we extract a total revenue of s.

Conversely, assume a solution to the pricing problem with total revenue s. Since only one copy of any item is available, the set of winning bids must define an independent set in G. As the maximum revenue from any bidder's bid is 1, there exists an independent set of size s in G. $\square$

4 Highway Problem with Monotonicity Constraint

A particularly intriguing special case of the single item prices problem is the 'highway problem' as introduced by Guruswami et al. [9]. Here, the items are edges of a simple path, and the bundles corresponding to bids requested by bidders are subpaths. The problem is NP-hard [4] by a simple transformation from Partition, and a $\log(m)$-approximation exists [2]. In this setting, it is natural to assume that the monotonicity constraint holds for any two *subpaths* only, but not necessarily for arbitrary subsets of items. For this problem we obtain the following results.

Theorem 5. *The highway problem with monotonicity constraint is NP-hard.*

The proof of this theorem is deferred to the full version of this paper. Notice that we cannot apply the PTAS from Theorem 2, as this crucially requires the monotonicity constraint for arbitrary subsets of items. Nevertheless, we derive an $O(\log B)$-approximation algorithm for the highway pricing problem with monotonicity constraint, where $B = \max_j b_j$. To this end, we present approximation guarantees for two special cases first.

Lemma 2. *The highway pricing problem with monotonicity constraint in which all bundles have size at least two is approximable within a factor of* 3 *by optimal uniform pricing.*

Proof. Consider an optimal solution with revenue OPT and let $p^*_{\max}$ be the highest item price in this solution. We claim that pricing all items at $p^*_{\max}/3$, yields a revenue of at least OPT/3. Clearly, an optimal uniform pricing is at least as good as the uniform $p^*_{\max}/3$ pricing.

First, we show that any winning bid $j \in W$ in the optimal pricing remains a winning bid for the uniform pricing at level $p^*_{\max}/3$. Let $|I_j| = \ell$. Then the valuation for bid j is at least $b_j \geq \lfloor \ell/2 \rfloor p^*_{\max}$, as by the monotonicity constraint the total price of any two consecutive items in an optimal solution is at least $p^*_{\max}$ and the bidder who placed bid j can afford the corresponding bundle I_j. In the uniform $p^*_{\max}/3$ pricing, the total bundle price is $\ell p^*_{\max}/3$, which is at most $\lfloor \ell/2 \rfloor p^*_{\max}$, for $\ell \geq 2$. In an optimal pricing, bundle I_j corresponding to bid j is priced at most at $\ell p^*_{\max}$, whereas in our uniform pricing, we get $\ell p^*_{\max}/3$. Hence, pricing all items $p^*_{\max}/3$ yields a revenue of at least OPT/3. □

The above lemma shows that whenever all bundles contain at least two items, we have a constant approximation. Now, we consider only instances in which bundles consist of exactly one item. Moreover, we restrict ourselves to instances in which $b_j/b_k \leq 2$, for any two bids j and k.

Property 1. Consider the highway pricing problem with monotonicity constraint, restricted to instances in which each bid is on a bundle with exactly one item and for any two bids j, k, it holds true that $b_j/b_k \leq 2$. Pricing each item at $\min_j b_j$ yields a revenue of at least OPT/2.

Theorem 6. *The best uniform pricing yields a solution with revenue at least* $\text{OPT}/(3+2\log_2 B)$ *for the highway pricing problem with monotonicity constraint, where* $B = \max_j b_j$. *Moreover, the time needed to find this solution is* $O(n^2 m)$.

Proof. Consider an optimum solution satisfying the monotonicity constraint, and let OPT_L denote the revenue of bidders whose bids are on bundles of size at least two and let OPT_r denote the revenue of bidders whose bids are on bundles of size one with valuation $2^{r-1} \le b_j < 2^r$ $(r = 1, \ldots, \lceil \log_2 B \rceil + 1)$ in this solution. Then $\text{OPT} = \text{OPT}_L + \sum_r \text{OPT}_r$.

Moreover, let APP_L denote the revenue obtained by the best uniform pricing and APP_r denote the revenue obtained by the best uniform pricing strategy for the bidders with bids in $J_r = \{j \in J : |I_j| = 1 \text{ and } 2^{r-1} \le b_j < 2^r\}$. By Property 1, we have that $\text{APP}_r \ge \text{OPT}_r/2$ and thus $\max_r \text{APP}_r \ge \sum_r \text{OPT}_r/(2\log_2 B)$. Moreover, from Lemma 2, it follows that $\text{APP}_L \ge \text{OPT}_L/3$. Hence, the solution found yields a revenue of

$$\max\{\text{APP}_L, \text{APP}_r : r = 1, \ldots, \lceil \log_2 B \rceil + 1\} \ge \text{OPT}/(3 + 2\log_2 B).$$

To see the claim on the time complexity, note that to find sufficiently good uniform pricing, we need to consider at most $O(n)$ different prices. For each price, we need to compute the set of winning bids and the revenue obtained on this price, which can be done in $O(nm)$ time. So, the best uniform price can be computed in $O(n^2 m)$ time. □

5 Conclusion

This paper studies purely algorithmic, or *omniscient* pricing problems, reflected by the fact that we assume bidders' valuations b_j to be known. Even more challenging are problems where valuations are private information, and incentive-compatible mechanisms are sought, that is, mechanisms that induce bidders to truthfully report their valuations. To that end, one might ask if previous ideas on the design of (random sampling) mechanisms, e.g. by Goldberg et al. [7] or Balcan et al. [3], can be applied. Particularly the latter paper suggests a general approach for reducing incentive-compatible mechanism design problems to the underlying algorithmic pricing problems. Among them, combinatorial auctions with single-minded bidders (the model considered in this paper). We leave this on the agenda for future research; complications might lay in the monotonicity constraint that we impose on the pricing function, respectively the general form of the affine price functions.

Acknowledgements

We thank the anonymous referees for some inspiring and helpful remarks.

References

1. Aggarwal, G., Feder, T., Motwani, R., Zhu, A.: Algorithms for multi-product pricing. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 72–83. Springer, Heidelberg (2004)
2. Balcan, M.F., Blum, A.: Approximation algorithms and online mechanisms for item pricing. In: Proc. 7th ACM Conf. Electronic Commerce, pp. 29–35. ACM, New York (2006)
3. Balcan, M.F., Blum, A., Hartline, J.D., Mansour, Y.: Mechanism design via machine learning. In: Proc. 46th IEEE Found. Comp. Sci., FOCS 2005, pp. 605–614 (2005)
4. Briest, P., Krysta, P.: Single-minded unlimited supply pricing on sparse instances. In: Proc. 17th ACM-SIAM Symp. Discr. Alg., SODA 2006, pp. 1093–1102 (2006)
5. Demaine, E.D., Feige, U., Hajiaghayi, M.T., Salavatipour, M.R.: Combination can be hard: Approximability of the unique coverage problem. In: Proc. 17th ACM-SIAM Symp. Discr. Alg., SODA 2006, pp. 162–171 (2006)
6. Garey, M.R., Johnson, D.S.: Computers and intractability: A guide to the theory of np-completeness. W. H. Freeman, New York (1979)
7. Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M., Wright, A.: Competitive auctions. Games and Economic Behavior 55, 242–269 (2006)
8. Grigoriev, A., van Loon, J., Sitters, R., Uetz, M.: How to sell a graph: Guidelines for graph retailers. In: Fomin, F.V. (ed.) WG 2006. LNCS, vol. 4271, pp. 125–136. Springer, Heidelberg (2006)
9. Guruswami, V., Hartline, J.D., Karlin, A.R., Kempe, D., Kenyon, C., McSherry, F.: On profit-maximizing envy-free pricing. In: Proc. 16th ACM-SIAM Symp. Discr. Alg., SODA 2005, pp. 1164–1173 (2005)
10. Hartline, J.D., Koltun, V.: Near-optimal pricing in near-linear time. In: Dehne, F., López-Ortiz, A., Sack, J.-R. (eds.) WADS 2005. LNCS, vol. 3608, pp. 422–431. Springer, Heidelberg (2005)
11. Walras, L.: Elements of pure economics. Allen and Unwin, London (1954)
12. Zuckerman, D.: Linear degree extractors and the inapproximability of max clique and chromatic number, ECCC Report TR05-100 (2005), `http://www.eccc.uni-trier.de/eccc/`